



Free Questions for Professional-Machine-Learning-Engineer

Shared by Rice on 29-01-2024

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)



Question 1

Question Type: MultipleChoice

You are training a deep learning model for semantic image segmentation with reduced training time. While using a Deep Learning VM Image, you receive the following error: The resource 'projects/deeplearning-platform/zones/europe-west4-c/acceleratorTypes/nvidia-tesla-k80' was not found. What should you do?

Options:

- A- Ensure that you have GPU quota in the selected region.
- B- Ensure that the required GPU is available in the selected region.
- C- Ensure that you have preemptible GPU quota in the selected region.
- D- Ensure that the selected GPU has enough GPU memory for the workload.

Answer:

B

Explanation:

The error message indicates that the selected GPU type (nvidia-tesla-k80) is not available in the selected region (europe-west4-c). This can happen when the GPU type is not supported in the region, or when the GPU quota is exhausted in the region. To avoid this error, you should ensure that the required GPU is available in the selected region before creating a Deep Learning VM Image. You can use the following steps to check the GPU availability and quota:

To check the GPU availability, you can use the `gcloud compute accelerator-types list` command with the `--filter` flag to specify the GPU type and the region. For example, to check the availability of nvidia-tesla-k80 in europe-west4-c, you can run:

```
gcloud compute accelerator-types list --filter'name=nvidia-tesla-k80 AND zone:europe-west4-c'
```

If the command returns an empty result, it means that the GPU type is not supported in the region. You can either choose a different GPU type or a different region that supports the GPU type. You can use the same command without the `--filter` flag to list all the available GPU types and regions. For example, to list all the available GPU types in europe-west4-c, you can run:

```
gcloud compute accelerator-types list --filter'zone:europe-west4-c'
```

To check the GPU quota, you can use the `gcloud compute regions describe` command with the `--format` flag to specify the region and the quota metric. For example, to check the quota for nvidia-tesla-k80 in europe-west4-c, you can run:

```
gcloud compute regions describe europe-west4-c --format'value(quotas.NVIDIA_K80_GPUS)'
```

If the command returns a value of 0, it means that the GPU quota is exhausted in the region. You can either request more quota from Google Cloud or choose a different region that has enough quota for the GPU type.

[Troubleshooting | Deep Learning VM Images | Google Cloud](#)

[Checking GPU availability](#)

[Checking GPU quota](#)

Question 2

Question Type: MultipleChoice

You recently trained an XGBoost model on tabular data. You plan to expose the model for internal use as an HTTP microservice. After deployment, you expect a small number of incoming requests. You want to productionize the model with the least amount of effort and latency. What should you do?

Options:

- A- Deploy the model to BigQuery ML by using CREATE MODEL with the BOOSTED-TREE-REGRESSOR statement and invoke the BigQuery API from the microservice.
- B- Build a Flask-based app. Package the app in a custom container on Vertex AI and deploy it to Vertex AI Endpoints.
- C- Build a Flask-based app. Package the app in a Docker image and deploy it to Google Kubernetes Engine in Autopilot mode.
- D- Use a prebuilt XGBoost Vertex container to create a model and deploy it to Vertex AI Endpoints.

Answer:

D

Explanation:

XGBoost is a popular open-source library that provides a scalable and efficient implementation of gradient boosted trees. You can use XGBoost to train a classification or regression model on tabular data. You can also use Vertex AI to productionize the model and expose it for internal use as an HTTP microservice. Vertex AI is a service that allows you to create and train ML models.

using Google Cloud technologies. You can use a prebuilt XGBoost Vertex container to create a model and deploy it to Vertex AI Endpoints. A prebuilt Vertex container is a container image that contains the dependencies and libraries needed to run a specific ML framework, such as XGBoost. You can use a prebuilt Vertex container to simplify the model creation and deployment process, without having to build your own custom container. Vertex AI Endpoints is a service that allows you to serve your ML models online and scale them automatically. You can use Vertex AI Endpoints to deploy the model from the prebuilt Vertex container and expose it as an HTTP microservice. You can also configure the endpoint to handle a small number of incoming requests, and optimize the latency and cost of serving the model. By using a prebuilt XGBoost Vertex container and Vertex AI Endpoints, you can productionize the model with the least amount of effort and latency. Reference:

[XGBoost documentation](#)

[Vertex AI documentation](#)

[Prebuilt Vertex container documentation](#)

[Vertex AI Endpoints documentation](#)

[Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate](#)

Question 3

Question Type: MultipleChoice

You recently developed a wide and deep model in TensorFlow. You generated training datasets using a SQL script that preprocessed raw data in BigQuery by performing instance-level transformations of the data.

a. You need to create a training pipeline to retrain the model on a weekly basis. The trained model will be used to generate daily recommendations. You want to minimize model development and training time. How should you develop the training pipeline?

Options:

A- Use the Kubeflow Pipelines SDK to implement the pipeline. Use the BigQueryJob component to run the preprocessing script and the customTrainingJob component to launch a Vertex AI training job.

B- Use the Kubeflow Pipelines SDK to implement the pipeline. Use the dataflowpythonjob component to preprocess the data and the customTraining Job component to launch a Vertex AI training job.

C- Use the TensorFlow Extended SDK to implement the pipeline. Use the Examplegen component

with the BigQuery executor to ingest the data the Transform component to preprocess the data, and the Trainer component to launch a Vertex AI training job.

D- Use the TensorFlow Extended SDK to implement the pipeline Implement the preprocessing steps as part of the input_fn of the model Use the ExampleGen component with the BigQuery executor to ingest the data and the Trainer component to launch a Vertex AI training job.

Answer:

C

Explanation:

TensorFlow Extended (TFX) is a platform for building end-to-end machine learning pipelines using TensorFlow. TFX provides a set of components that can be orchestrated using either the TFX SDK or Kubeflow Pipelines. TFX components can handle different aspects of the pipeline, such as data ingestion, data validation, data transformation, model training, model evaluation, model serving, and more. TFX components can also leverage other Google Cloud services, such as BigQuery, Dataflow, and Vertex AI.

Why not A: Using the Kubeflow Pipelines SDK to implement the pipeline is a valid option, but using the BigQueryJobOp component to run the preprocessing script is not optimal. This would require writing and maintaining a separate SQL script for data transformation, which could introduce inconsistencies and errors. It would also make it harder to reuse the same preprocessing logic for both training and serving.

Why not B: Using the Kubeflow Pipelines SDK to implement the pipeline is a valid option, but using the DataflowPythonJobOp component to preprocess the data is not optimal. This would require writing and maintaining a separate Python script for data transformation, which could introduce inconsistencies and errors. It would also make it harder to reuse the same preprocessing logic for both training and serving.

Why not D: Using the TensorFlow Extended SDK to implement the pipeline is a valid option, but implementing the preprocessing steps as part of the input_fn of the model is not optimal. This would make the preprocessing logic tightly coupled with the model code, which could reduce modularity and flexibility. It would also make it harder to reuse the same preprocessing logic for both training and serving.

Question 4

Question Type: MultipleChoice

You are building a predictive maintenance model to preemptively detect part defects in bridges. You plan to use high definition images of the bridges as model inputs. You need to explain the

output of the model to the relevant stakeholders so they can take appropriate action. How should you build the model?

Options:

- A- Use scikit-learn to build a tree-based model, and use SHAP values to explain the model output.
- B- Use scikit-learn to build a tree-based model, and use partial dependence plots (PDP) to explain the model output.
- C- Use TensorFlow to create a deep learning-based model and use Integrated Gradients to explain the model output.
- D- Use TensorFlow to create a deep learning-based model and use the sampled Shapley method to explain the model output.

Answer:

C

Explanation:

According to the official exam guide¹, one of the skills assessed in the exam is to "explain the predictions of a trained model". TensorFlow² is an open source framework for developing and deploying machine learning and deep learning models. TensorFlow supports various model explainability methods, such as Integrated Gradients³, which is a technique that assigns an importance score to each input feature by approximating the integral of the gradients along the path from a baseline input to the actual input. Integrated Gradients can help explain the output of a deep learning-based model by highlighting the most influential features in the input images. Therefore, option C is the best way to build the model for the given use case. The other options are not relevant or optimal for this scenario. Reference:

[Professional ML Engineer Exam Guide](#)

[TensorFlow](#)

[Integrated Gradients](#)

[Google Professional Machine Learning Certification Exam 2023](#)

[Latest Google Professional Machine Learning Engineer Actual Free Exam Questions](#)

Question 5

Question Type: MultipleChoice

You are developing a model to predict whether a failure will occur in a critical machine part. You have a dataset consisting of a multivariate time series and labels indicating whether the machine part failed. You recently started experimenting with a few different preprocessing and modeling approaches in a Vertex AI Workbench notebook. You want to log data and track artifacts from each run. How should you set up your experiments?

A.

1. Use the Vertex AI SDK to create an experiment, and set up Vertex ML Metadata.
2. Use the `log_time_series_metrics` function to track the preprocessed data, and use the `log_metrics` function to log loss values.

B.

1. Use the Vertex AI SDK to create an experiment, and set up Vertex ML Metadata.
2. Use the `log_time_series_metrics` function to track the preprocessed data, and use the `log_metrics` function to log loss values.

C.

1. Create a Vertex AI TensorBoard instance, and use the Vertex AI SDK to create an experiment and associate the TensorBoard instance.
2. Use the `assign_input_artifact` method to track the preprocessed data, and use the `log_time_series_metrics` function to log loss values.

D.

1. Create a Vertex AI TensorBoard instance, and use the Vertex AI SDK to create an experiment and associate the TensorBoard instance.
2. Use the `log_time_series_metrics` function to track the preprocessed data, and use the `log_metrics` function to log loss values.

Options:

- A- Option A
- B- Option B
- C- Option C
- D- Option D

Answer:

A

Explanation:

The option A is the most suitable solution for logging data and tracking artifacts from each run of a model development experiment in a Vertex AI Workbench notebook. Vertex AI Workbench is a service that allows you to create and run interactive notebooks on Google Cloud. You can use Vertex AI Workbench to experiment with different preprocessing and modeling approaches for your time series prediction problem. You can also use the Vertex AI TensorBoard instance and the Vertex AI SDK to create an experiment and associate the TensorBoard instance. TensorBoard is a tool that allows you to visualize and monitor the metrics and artifacts of your ML experiments. You can use the Vertex AI SDK to create an experiment object, which is a logical grouping of runs that share a common objective. You can also use the Vertex AI SDK to associate the experiment object with a TensorBoard instance, which is a managed service that hosts a TensorBoard web app. By using the Vertex AI TensorBoard instance and the Vertex AI SDK, you can easily set up and manage your experiments, and access the TensorBoard web app from the Vertex AI console. You can also use the `log_time_series_metrics` function and the `log_metrics` function to log data and track artifacts from each run. The `log_time_series_metrics` function is a function that allows you to log the time series data, such as the multivariate time series and the labels, to the TensorBoard instance. The `log_metrics` function is a function that allows you to log the scalar metrics, such as the loss values, to the TensorBoard instance. By using these functions, you can record the data and artifacts from each run of your experiment, and compare them in the TensorBoard web app. You can also use the TensorBoard web app to visualize the data and artifacts, such as the time series plots, the scalar charts, the histograms, and the distributions. By using the Vertex AI TensorBoard instance, the Vertex AI SDK, and the log functions, you can log data and track artifacts from each run of your experiment in a Vertex AI Workbench notebook. Reference:

[Vertex AI Workbench documentation](#)

[Vertex AI TensorBoard documentation](#)

[Vertex AI SDK documentation](#)

[log_time_series_metrics function documentation](#)

[log_metrics function documentation](#)

[Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate]

Question 6

Question Type: MultipleChoice

Your organization wants to make its internal shuttle service route more efficient. The shuttles currently stop at all pick-up points across the city every 30 minutes between 7 am and 10 am. The development team has already built an application on Google Kubernetes Engine that

requires users to confirm their presence and shuttle station one day in advance. What approach should you take?

Options:

- A-** 1. Build a tree-based regression model that predicts how many passengers will be picked up at each shuttle station.
2. Dispatch an appropriately sized shuttle and provide the map with the required stops based on the prediction.
- B-** 1. Build a tree-based classification model that predicts whether the shuttle should pick up passengers at each shuttle station.
2. Dispatch an available shuttle and provide the map with the required stops based on the prediction
- C-** 1. Define the optimal route as the shortest route that passes by all shuttle stations with confirmed attendance at the given time under capacity constraints.
2 Dispatch an appropriately sized shuttle and indicate the required stops on the map
- D-** 1. Build a reinforcement learning model with tree-based classification models that predict the presence of passengers at shuttle stops as agents and a reward function around a distance-based metric
2. Dispatch an appropriately sized shuttle and provide the map with the required stops based on the simulated outcome.

Answer:

A

Explanation:

This answer is correct because it uses a regression model to estimate the number of passengers at each shuttle station, which is a continuous variable. A tree-based regression model can handle both numerical and categorical features, such as the time of day, the location of the station, and the weather conditions. Based on the predicted number of passengers, the organization can dispatch a shuttle that has enough capacity and provide a map that shows the required stops. This way, the organization can optimize the shuttle service route and reduce the waiting time and fuel consumption. Reference:

[Tree-based regression models]

To Get Premium Files for Professional-Machine-Learning-Engineer Visit

<https://www.p2pexams.com/products/professional-machine-learning-engineer>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-machine-learning-engineer>

20%
DISCOUNT

P2P
exams